



IBM POWER AC922

The Best Server for Enterprise AI

As AI initiatives shift from backroom experiments to boardroom imperatives, many organizations find their existing infrastructure ill-suited to help them make the journey. IT leaders are now under pressure to implement untested tools, circumvent proven processes, and shrink the timeframe from beta to production. Organizations that do not evaluate their existing IT infrastructure, ensuring it is appropriate for Enterprise AI, will risk getting caught in an endless loop of failed experiments.

The IBM Power System AC922 supports Enterprise AI initiatives throughout their lifecycle. With an optimized hardware and software stack, the AC922 delivers earlier prototypes and yields faster ROI. The world's only server enabling NVIDIA NVLink between CPUs and GPUs, the Power AC922 delivers 5.6x the data movement of x86-based servers. This I/O enables GPU-based algorithms to leverage system memory, enabling models up to 60x larger than x86 based systems. The Power AC922 delivers faster AI insights, providing up to 3.8x5 speedups for deep learning workloads. Proven as the backbone of the world's largest supercomputers, the Power AC922 is capable of meeting an enterprise's loftiest AI aspirations.

KEY FEATURES

- More accuracy - GPUs access system RAM for larger models
- Faster insights - significant deep learning speedups
- Rapid deployment - integrated software & hardware stack
- Enterprise-ready- cutting edge AI IBM Power reliability & IBM support
- Proven at scale - from one node to the world's fastest supercomputers

ACCELERATED COMPUTE

Yields more accuracy (up to 60x larger models: uniquely capable of supporting larger models and data sets, by accessing system from GPU-based processes and algorithms, without PCIe bottlenecks in x86-based systems.

Delivers faster insights: up to 3.8x5 the AI performance, vs similarly-configured x86-based systems.

Provides faster I/O: includes a variety of next-generation I/O architectures, including: PCIe Gen4, CAPI



PCIe Gen4 provides 2x data throughput vs. PCIe gen 3

Simplest AI Architecture-Coherence simplifies coding by abstracting data movement and locality for developers

5.6x more I/O bandwidth

Enables faster insights -3.8x speedup based on competitor solutions

ABOUT BIOS IT

BIOS IT is a global design house, systems builder and integrated solution provider for enterprise performance computing. We construct bespoke clusters and appliances by hand picking the best components and newest technologies based on specific customer goals. We support well-known organisations at the top of their research fields, across science, engineering, academia and finance disciplines. Our comprehensive range of products and services include: high-performance enterprise servers, storage and networking, on-premise or in the cloud, with associated services, support, hosting and software.

2.0, OpenCAPI and NVIDIA NVLINK.

CPUs unleash GPU acceleration: supports up to 5.6x1 more I/O and 2x more threads than its x86 contemporaries - configurations with 16, 18, 20 and 22 cores, for up to 44 cores.

Includes the most advanced data center GPUs: pairs POWER9 CPUs and NVIDIA Tesla V100 with NVLink GPUs delivering up to 5.6x1 times the performance for each pairing.

1st PCIe Gen4 server: PCIe Gen4 delivers approximately 2x the data bandwidth² of the PCIe Gen3 interconnect found in x86-based servers.

Simplest AI architecture: By leveraging the 5.6x1 faster NVIDIA NVLink interconnect, sharing memory between CPUs and GPUs doesn't bottleneck down to PCIe Gen3 speeds, as it would on x86-based servers.

Enterprise-ready: start with one node, then scale to a rack or thousands of nodes with near linear scaling efficiency.

	BIGTWIN™
MICROPROCESSORS	2x POWER9 with NVLink CPUs 16, 20 cores, or 18, 22 cores (with water cooling)
L2 CACHE	512 K
L3 CACHE	10 MB
MEMORY	Up to 2 TB, from 16 DDR4 RDIMM Sockets
STORAGE	2X SFF (2.5") DRIVE BAYS, OPTIONAL NVME SSD SUPPORT IN PCIE SLOTS
ADAPTOR SLOTS	4 or 6 SXM 2.0 sockets, for NVIDIA Tesla V100 GPU Accelerators with NVLink. 2x PCIe x16 4.0 slots 1x PCIe16x (x8,x8) 4.0 slot (multi-socket host direct supported) 1x PCIe x4 4.0 slot
OS	Red Hat Enterprise Linux, Ubuntu Linux

WHY IBM?

IBM is a trailblazer in AI—From early machine learning system in IBM Research to Watson on Jeopardy, AI isn't just a buzzword for IBM. And we're applying that innovation to cognitive infrastructure, helping our customers on their journey to AI.

For more information visit www.bios-it.com.



www.bios-it.com
AMERICAS: 1-800-654-BIOS | EMEA: +44 (0) 203 178 6467 | APAC: +61 (0)2 8866 3343

All products and companies referred to herein are trademarks or registered trademarks of their respective companies or mark holders.